

MICROSOFT CERTIFIED: AZURE DATA FUNDAMENTALS (DP-900) PRACTICE TEST (SAMPLE) STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

[https://
mscertifiedazuredatafundamentals.examzify.com](https://mscertifiedazuredatafundamentals.examzify.com)



Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	16

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 – 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

- 1. What distinguishes a clustered index from a regular index?**
 - A. It requires a copy of the indexed data
 - B. It sorts the actual table data based on the index
 - C. It improves data security and encryption
 - D. It allows for unlimited data storage capacity
- 2. Which formats are considered types of batch data?**
 - A. CSV, JSON, and Apache Kafka
 - B. Only CSV and TSV
 - C. JSON, XML, and Blob files
 - D. CSV, TSV, JSON, XML, Parquet, Blob files, and others
- 3. What functionality does Polybase provide in Azure Synapse Analytics?**
 - A. It processes SQL queries only
 - B. It retrieves data from relational and non-relational sources
 - C. It mainly serves as a data visualization tool
 - D. It directly interfaces with Azure Cosmos DB
- 4. What is the purpose of an App in Power BI?**
 - A. A collection of preset visuals and reports shared with an organization
 - B. A tool for data cleansing and preparation
 - C. A method for exporting reports to PDF
 - D. A feature for importing data from external sources
- 5. What latency does Cosmos DB guarantee for reads and writes at the 99th percentile globally?**
 - A. Less than 5 ms
 - B. Less than 10 ms
 - C. Less than 20 ms
 - D. Less than 50 ms
- 6. Which of the following tools can be used to provision services in Azure?**
 - A. Only Azure Portal
 - B. Azure Portal and Azure CLI
 - C. Azure CLI and Azure Powershell
 - D. All of the above

- 7. Azure Synapse Analytics is designed to quickly process what type of data?**
- A. Small local data files
 - B. Large amounts of data
 - C. Medium-sized data sets
 - D. Transactional data only
- 8. Which of the following is a characteristic of a report in Power BI?**
- A. It is a single visualization
 - B. It is a collection of visualizations spanning one or more pages
 - C. It interacts directly with users
 - D. It cannot be edited after publication
- 9. Which of the following describes Graph databases in NoSQL?**
- A. They focus on data stored in a tabular format
 - B. They model relationships between entities efficiently
 - C. They do not support indexing
 - D. They are limited to structured data only
- 10. When should you consider using Polybase?**
- A. When using Azure Blob Storage
 - B. When ingesting data with Azure Synapse
 - C. When working with SQL Server
 - D. When creating Power BI reports

Answers

SAMPLE

1. B
2. D
3. B
4. A
5. B
6. D
7. B
8. B
9. B
10. B

SAMPLE

Explanations

SAMPLE

1. What distinguishes a clustered index from a regular index?

- A. It requires a copy of the indexed data
- B. It sorts the actual table data based on the index**
- C. It improves data security and encryption
- D. It allows for unlimited data storage capacity

A clustered index is distinct from a regular index primarily because it dictates the physical order of the data within the table itself. When a clustered index is created on a table, the rows of the table are sorted and stored on disk according to the indexed columns. This means that the data is organized in the same order as defined by the clustered index, which allows for efficient retrieval of data, especially in queries that sort or filter by the indexed columns. Because the clustered index determines how the data is stored, there can only be one clustered index per table, as the data can only be physically arranged in one order at a time. This feature of a clustered index contrasts with a non-clustered index, which maintains a separate structure that references the actual table rows and does not affect their physical order. The other options address characteristics that aren't core distinctions of clustered versus regular indexes. For instance, a copy of the indexed data might be relevant in specific contexts, but it does not define the fundamental difference. Data security and encryption relate to database protection measures, while unlimited data storage capacity is not a feature of either index type. Thus, the correct understanding of a clustered index is focused on its role in sorting and storing the actual data in the table in a

2. Which formats are considered types of batch data?

- A. CSV, JSON, and Apache Kafka
- B. Only CSV and TSV
- C. JSON, XML, and Blob files
- D. CSV, TSV, JSON, XML, Parquet, Blob files, and others**

Batch data refers to a collection of data that is processed in groups or "batches" rather than individually and is typically characterized by being stored and processed in bulk. The correct choice identifies a comprehensive list of formats that are commonly used for batch data storage. CSV (Comma-Separated Values) and TSV (Tab-Separated Values) are well-known formats for representing tabular data and are frequently utilized for batch data operations due to their simplicity and wide compatibility. JSON (JavaScript Object Notation) and XML (eXtensible Markup Language) are also prevalent formats for structured data, allowing for hierarchical representation, which makes them suitable for batch processing. Parquet is a columnar storage file format that is optimized for batch processing and can significantly improve performance in data analytics scenarios. Blob files, particularly in the context of cloud storage, can encapsulate various types of data, including images, documents, and other file types, making them suitable for batch processing of large datasets. By including all these formats—CSV, TSV, JSON, XML, Parquet, and Blob files—the selected choice encompasses a wide variety of batch data formats that are utilized across different platforms and applications, reflecting the diverse nature of batch data processing in the field of data management.

3. What functionality does Polybase provide in Azure Synapse Analytics?

- A. It processes SQL queries only
- B. It retrieves data from relational and non-relational sources**
- C. It mainly serves as a data visualization tool
- D. It directly interfaces with Azure Cosmos DB

PolyBase in Azure Synapse Analytics is designed to enable you to retrieve data from various data sources, including both relational and non-relational sources. This functionality allows users to efficiently query large datasets that reside outside of the traditional relational database system. By using PolyBase, you can integrate data from sources such as Azure Blob Storage or Hadoop into your analytical workflows, leveraging the power of Azure Synapse without needing to physically move or replicate that data into the Synapse environment. This capability makes PolyBase particularly valuable for handling big data scenarios, as it simplifies the process of analyzing data that is stored in different formats or systems. It allows users to write T-SQL queries that can seamlessly access and combine data from a variety of sources, enhancing the flexibility and power of your analytical capabilities in Azure Synapse Analytics.

4. What is the purpose of an App in Power BI?

- A. A collection of preset visuals and reports shared with an organization**
- B. A tool for data cleansing and preparation
- C. A method for exporting reports to PDF
- D. A feature for importing data from external sources

The purpose of an App in Power BI is to serve as a collection of preset visuals and reports that are shared within an organization. This allows users to create a cohesive and organized experience for accessing data insights, tailored for specific business needs or functions. When an App is created in Power BI, it typically bundles various reports, dashboards, and datasets that have been curated for a particular audience. This enables users to interact with the data they need without having to navigate through multiple reports or dashboards individually. It promotes collaboration and ensures that relevant information is readily accessible to stakeholders, thereby enhancing decision-making processes. The other options, while they describe useful features or functionalities within the Power BI ecosystem, do not accurately capture the essence of what an App is designed to do. For example, data cleansing and preparation are important aspects of data management, but they are handled by tools such as Power Query, not Apps. Similarly, exporting reports to PDF or importing data from external sources are distinct functionalities that do not define the purpose of an App in Power BI.

5. What latency does Cosmos DB guarantee for reads and writes at the 99th percentile globally?

- A. Less than 5 ms
- B. Less than 10 ms**
- C. Less than 20 ms
- D. Less than 50 ms

Cosmos DB is designed to provide low-latency performance, which is essential for global applications requiring fast access to data. The service guarantees that for reads and writes, the latency will be less than 10 milliseconds at the 99th percentile globally. This means that in 99% of the cases, read and write operations will complete within this time frame, offering consistent and high-performance data operations even when distributed across various geographical locations. This guarantee is a key feature of Azure Cosmos DB, which is highly regarded for its scalability, availability, and responsiveness. By ensuring that such performance metrics are met, Azure Cosmos DB supports applications that demand real-time processing, enhancing the user experience and ensuring that data-driven decisions can be made promptly. As a result, businesses can trust Cosmos DB to handle their workloads efficiently while maintaining a high standard of data access performance.

6. Which of the following tools can be used to provision services in Azure?

- A. Only Azure Portal
- B. Azure Portal and Azure CLI
- C. Azure CLI and Azure Powershell
- D. All of the above**

All the mentioned tools are valid options for provisioning services in Azure, which is why the correct choice includes them all. The Azure Portal is a web-based interface that allows users to manage and provision Azure resources visually. It is user-friendly and provides a graphical interface, making it accessible for users who prefer not to use command-line tools. Azure CLI is a command-line interface that enables users to manage Azure resources through commands in a shell. It is particularly useful for automation and scripting, allowing users to interact with Azure services programmatically. Azure PowerShell is another command-line utility that provides a scripting environment specifically for Windows users. It allows for handling Azure resources through PowerShell commands, which can be beneficial for those who are already familiar with the PowerShell environment. Since each of these tools effectively allows for the provisioning of services within Azure, the correct answer encompasses all options as they each serve distinct purposes and cater to different user preferences and requirements in managing Azure resources.

7. Azure Synapse Analytics is designed to quickly process what type of data?

- A. Small local data files
- B. Large amounts of data**
- C. Medium-sized data sets
- D. Transactional data only

Azure Synapse Analytics is specifically designed to handle and quickly process large amounts of data, making it particularly suitable for big data analytics and complex queries. It integrates big data and data warehousing capabilities, allowing users to analyze both structured and unstructured data efficiently. This is essential for organizations that need to process and derive insights from vast datasets, such as those found in data lakes or extensive enterprise databases. The platform utilizes massively parallel processing (MPP) architecture, which enables it to distribute workloads across many nodes, resulting in enhanced performance for aggregating, sorting, and querying vast quantities of data simultaneously. This capability makes Azure Synapse Analytics a powerful tool for business intelligence, analytics, and reporting, especially when dealing with large-scale data scenarios typical in modern data environments. Regarding the other options, while Azure Synapse can technically handle medium-sized data sets or smaller files, its primary strengths and optimizations are geared towards large amounts of data, making it overqualified for smaller data processing tasks. Transactional data can also be managed, but this is not the exclusive focus of Azure Synapse Analytics. Its design is broader, aimed at facilitating comprehensive analytics across various data types at significant scales.

8. Which of the following is a characteristic of a report in Power BI?

- A. It is a single visualization
- B. It is a collection of visualizations spanning one or more pages**
- C. It interacts directly with users
- D. It cannot be edited after publication

A report in Power BI is best described as a collection of visualizations that can span one or more pages. This characteristic allows users to present related data insights in a cohesive manner, providing the flexibility to include various types of visualizations such as charts, tables, and maps, all designed to display the insights derived from data. The fact that reports can consist of multiple pages enhances the capability to explore complex datasets by enabling different views and perspectives within a single report framework. Users can navigate between pages to examine different aspects of the data, making it a powerful tool for data analysis and presentation. Furthermore, the nature of a report as a collection of visualizations implies that it can convey a comprehensive story or analysis about the data, as opposed to being a single visualization that limits the narrative to one specific view.

9. Which of the following describes Graph databases in NoSQL?

- A. They focus on data stored in a tabular format
- B. They model relationships between entities efficiently**
- C. They do not support indexing
- D. They are limited to structured data only

Graph databases in NoSQL are designed specifically to efficiently model and manage relationships between entities. This is a fundamental characteristic that sets them apart from other types of databases, such as relational or document databases. In a graph database, data is represented as nodes (entities) and edges (relationships), which allows for complex queries that involve traversing relationships. This structure is particularly advantageous for scenarios such as social networks, recommendation systems, and network analysis, where the relationships among data points are as important as the data itself. The other options highlight aspects that are either unrelated to the core functionality of graph databases or incorrectly characterize them. For instance, graph databases do not utilize a tabular format; instead, they focus on interconnected data, making them distinct from traditional relational databases. Moreover, many graph databases do support indexing to optimize data retrieval, which contradicts the idea that they do not support indexing. Lastly, graph databases are not limited to structured data; they can handle semi-structured and unstructured data too, allowing flexibility in how data is modeled and queried.

10. When should you consider using Polybase?

- A. When using Azure Blob Storage
- B. When ingesting data with Azure Synapse**
- C. When working with SQL Server
- D. When creating Power BI reports

Polybase is a technology that allows you to query and import data from different data sources directly into a data warehouse without the need to move the data physically. This capability is particularly advantageous when you are working with large datasets. Considering when to use Polybase, it is most suitable when ingesting data with Azure Synapse Analytics. Azure Synapse is designed to integrate analytics and data warehousing, and Polybase facilitates the seamless querying of external data sources, including Azure Blob Storage, Azure Data Lake, and other SQL Server instances. This allows data engineers and analysts to combine data from both relational and non-relational sources easily during the data ingestion process, leading to more flexible and efficient data workflows. Other scenarios, while relevant to data management and analytics, do not leverage the unique capabilities of Polybase to the same extent. For example, while working with Azure Blob Storage, the data can be accessed directly, but Polybase is most effective when integrated within the broader capabilities of Azure Synapse, where data ingestion and hybrid queries are essential.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

<https://mscertifiedazuredatafundamentals.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE